




INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification ⁶ : H04Q 11/04		A1	(11) International Publication Number: WO 98/08355 (43) International Publication Date: 26 February 1998 (26.02.98)
(21) International Application Number: PCT/CA97/00610 (22) International Filing Date: 14 August 1997 (14.08.97) (30) Priority Data: 60/024,023 16 August 1996 (16.08.96) US		(81) Designated States: CA, JP, European patent (AT, BE, CH, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE). Published <i>With international search report.</i>	
(71) Applicant: NORTHERN TELECOM LIMITED [CA/CA]; World Trade Center of Montreal, 8th floor, 380 St. Antoine Street West, Montreal, Quebec H2Y 3Y4 (CA). (72) Inventor: VALLEE, Richard; 67 Rue Castillou, Gatineau, Quebec J8T 5Y4 (CA). (74) Agent: TOYOOKA, Yoshiharu; Northern Telecom Limited, Patent Dept., P.O. Box 3511, Station "C", Ottawa, Ontario K1Y 4H7 (CA).			

(54) Title: INVERSE MULTIPLEXING OF DIGITAL DATA



(57) Abstract

In ATM networks, digital data in ATM cells are sent to a destination node over more than one transmission link in round robin fashion. This is called inverse multiplexing. At connection start-up, the source node informs the destination node of the specific round robin fashion of the transmission links so that the ATM cells are reassembled in a proper sequential order. Inverse multiplexing control cells are used to communicate between the source node and destination node for connectivity testing of transmission links. Cell stuffing is also provided in one embodiment to accommodate non-synchronized links among transmission links. In a particular embodiment, two consecutive control cells indicate a stuffing cell. A start-up procedure is described when not all the transmission links are usable.

FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AL	Albania	ES	Spain	LS	Lesotho	SI	Slovenia
AM	Armenia	FI	Finland	LT	Lithuania	SK	Slovakia
AT	Austria	FR	France	LU	Luxembourg	SN	Senegal
AU	Australia	GA	Gabon	LV	Latvia	SZ	Swaziland
AZ	Azerbaijan	GB	United Kingdom	MC	Monaco	TD	Chad
BA	Bosnia and Herzegovina	GE	Georgia	MD	Republic of Moldova	TG	Togo
BB	Barbados	GH	Ghana	MG	Madagascar	TJ	Tajikistan
BE	Belgium	GN	Guinea	MK	The former Yugoslav Republic of Macedonia	TM	Turkmenistan
BF	Burkina Faso	GR	Greece	ML	Mali	TR	Turkey
BG	Bulgaria	HU	Hungary	MN	Mongolia	TT	Trinidad and Tobago
BJ	Benin	IE	Ireland	MR	Mauritania	UA	Ukraine
BR	Brazil	IL	Israel	MW	Malawi	UG	Uganda
BY	Belarus	IS	Iceland	MX	Mexico	US	United States of America
CA	Canada	IT	Italy	NZ	Niger	UZ	Uzbekistan
CF	Central African Republic	JP	Japan	NL	Netherlands	VN	Viet Nam
CG	Congo	KE	Kenya	NO	Norway	YU	Yugoslavia
CH	Switzerland	KG	Kyrgyzstan	NZ	New Zealand	ZW	Zimbabwe
CI	Côte d'Ivoire	KP	Democratic People's Republic of Korea	PL	Poland		
CM	Cameroon	KR	Republic of Korea	PT	Portugal		
CN	China	KZ	Kazakhstan	RO	Romania		
CU	Cuba	LC	Saint Lucia	RU	Russian Federation		
CZ	Czech Republic	LI	Liechtenstein	SD	Sudan		
DE	Germany	LK	Sri Lanka	SE	Sweden		
DK	Denmark	LR	Liberia	SG	Singapore		
EE	Estonia						

INVERSE MULTIPLEXING OF DIGITAL DATA

Field of the Invention

The invention relates to a new mechanism for sending ATM cells transparently over multiple slower transmission links. In particular, the invention is directed to a method of inverse multiplexing of a series of ATM cells transparently over N transmission links (N being a positive integer) of slower speed.

10 Background of the Invention

15 It has been recognized that the T1/E1 rate (1.544/2.048 Mbit/s) is a cost effective way of user access to an ATM network as well as connection between ATM network switches. However, as ATM technology for wide area networks is deployed more and more, demands for transmission links of a rate higher than T1/E1 are increasing. Links of higher rates, such as T3/E3 (44.736/34.368 Mbit/s), have been designed to meet these needs. However, the cost of T3/E3 links is still prohibitive in many cases and the ratio of cost versus realistic utilization of the entire rate is not always attractive and fully justified for new ATM end users and service providers. ATM inverse multiplexers (AIMs) have been proposed to satisfy the need by using multiple T1/E1 links which are grouped collectively to provide the service at a higher rate.

Figure 1 and Figure 2 show two sample configurations in which
25 AIMs are used. Figure 1 depicts a user access to a network through
user network interfaces (UNIs) and Figure 2 a link connection between
ATM switches through broadband inter-carrier interfaces (BICIs) or
private network to network interfaces (PNNIs).

Referring to the figures, the basic function of AIMs is to work in pairs to take an ATM cell stream coming from the ATM layer, send it over the multiple links by spreading cells over the available links and ensure that the initial cell stream can be retrieved at the far end. Thus the AIMs preferably make the ATM traffic transparent to the ATM layer over multiple links which connect them. As far as the ATM layer is concerned, it should only see a pipe whose rate is now the sum of the multiple link rates. It is assumed that each link is run in clear-mode without the presence of intermediate ATM nodes processing ATM

cells. This means that there should be no cell discard by any intermediate transmission equipment.

Currently no ATM inverse multiplexing protocols have been proposed which can properly interwork existing ATM inverse multiplexers or other ATM products which are already available on the market, and yet are flexible enough to fit into the current standard ATM specifications. Two proposals for an ATM inverse multiplexing protocol have so far been made and are described in detail below.

10 New Transmission Convergence Protocol Using GFC Bits

This protocol was presented in "Physical Layer Sub-Working Group ATM Forum/94-0775, ATM Inverse Multiplexing Mechanism", September 1994, by StrataCom Inc. The protocol robs two of the Generic Flow Control (GFC) bits contained in each cell transmitted over the 15 multiple T1/E1 links to implement a new transmission convergence (TC) layer. Figure 3 shows the ATM cell structure which is defined in the ITU Recommendation. The TC layer is defined by one GFC bit for framing and the other one for link control. The framing bit is used to determine relative link delays while the link control bit is used for 20 communication, control and administration between two TC points at two ends of the inverse multiplexer.

In order to establish the sequence of cells over the links in a round robin manner, one end is defined as being "master" and the 25 other as "slave". The "master" decides and informs the slave about the multiple link configuration using the control channel implemented through the link control bits.

This protocol is only applicable, however, for UNI application points because GFC bits that are robbed to implement the TC layer are only present in a cell defined for UNI. For NNI cells, the 30 corresponding bits are no longer available since they are captured under the VPI field. Service providers are interested in ATM inverse multiplexers for carrying ATM traffic at rates higher than T1/E1 and lower than T3/E3, but this protocol will not satisfy their need. It should also be noted that the protocol calls for a need to identify a "master" 35 and a "slave" TC point and that requires an additional setting to be performed by the network operator.

Bit Pipe Inverse Multiplexing

5 This protocol was presented in "Physical Layer Sub-Working Group ATM Forum/94-0956, Inverse Multiplexing of ATM cells over low speed UNIs such as T1 and E1", September 1994, by Digital Link Corporation. It proposes a "bit pipe" inverse multiplexing technique requiring the definition of a "bonding" (bandwidth on demand) like specification for N (positive number) T1/E1 inverse multiplexing.

10 It is not clear in the proposal how both ends of the links exchange information concerning the order of cells to be transferred from one end to another over multiple links. The proposal mentions the existence and deployment of physical layer protocols that perform inverse multiplexing. The inverse multiplexer which can be used in this proposal is presumably the one defined by Digital Link Corporation in their "DL3800 DS1 Inverse Multiplexer Users Manual, 1993".

15 The inverse multiplexing protocol defined in the above user's manual relies on the definition of an extra bit taken from T1/E1 payload bits to configure the multiple links and adjust differential link delays. This protocol introduces the need for extra processing of data between devices dealing with T1/E1 frames and ATM cell delineation. 20 It also causes the ATM cells to no longer be byte aligned with the DS1/E1 frame. This is a requirement by the ATM Forum UNI DS1/E1 Physical Layer specifications. Changes like this would not be welcome by end users, vendors and service providers who are already using and deploying ATM equipment.

25 U. S. Patent No. 5,608,733, Mar. 4, 1997, Vallee et al, describes good ways of obviating the above noted problems. The patent uses ATM sequence number cells indicating a specific round robin order of a plurality of transmission links over which ATM data cells are transmitted. The ATM sequence number cells also indicate whether or 30 not a destination is ready to receive ATM data cells in that specific round robin order.

The present invention extends further variety of functionalities which are useful in inverse multiplexing.

35 Objects of the Invention

It is therefore an object of the invention to provide a method of sending ATM traffic over a connection consisting of a plurality of transmission links.

5 It is another object of the invention to provide a method of sending a series of ATM cells spread over a plurality of transmission links in a specific round robin order.

It is yet another object of the invention to provide a method of sending ATM cells containing sequence numbers therein over a plurality of transmission links in a specific round robin order.

10 It is a further object of the invention to provide a method of sending a series of ATM cells between ATM inverse multiplexers over a connection consisting of a plurality of transmission links, transparent to the ATM layer.

15 It is yet another object of the invention to provide a method of sending ATM traffic over a connection which consists of a plurality of transmission links and has been reconfigured after the connection start-up.

20 It is still another object of the invention to provide an ATM inverse multiplexing method which is applicable to UNIs, BICIs and PNNs.

It is a further object of the invention to provide a method of preserving link integrity by periodically sending sequence number cells.

It is a further object of the invention to provide a method of handling link failure and link reconfiguration.

25 It is still a further object of the invention to provide a method of stuffing cells to accommodate non-synchronized links.

Summary of the Invention

Briefly stated, the invention reside in the inverse multiplexing 30 digital data over a connection consisting of a plurality of transmission links, said data containing a series of ATM data cells. According to one aspect, the invention is directed to a method of testing connectivity of the transmission links. The method comprises steps of sending a series of inverse multiplexing control cells containing a test message over 35 one of the transmission links, and receiving a series of inverse multiplexing control cells containing received test message over all the remaining transmission links. The method further includes a step of

verifying the received test message to determine the connectivity of the transmission links.

According to another aspect, the invention is directed to an apparatus for testing the connectivity of transmission links used in the 5 inverse multiplexing.

In accordance with a yet another aspect, the method further includes a step of sending two consecutive inverse multiplexing control cells in a frame, indicating cell stuffing.

According to a further aspect, the invention is directed to a 10 method of inverse multiplexing digital data over a connection consisting of a plurality of transmission links, said data containing a series of ATM data cells. In particular, the method comprises steps of identifying N number of transmission links over which the series of ATM data cells are to be transmitted, N being a positive integer and 15 sending a series of inverse multiplexing control cells indicating a first specific round robin order of the transmission links. The method further includes steps of receiving a series of inverse multiplexing control cells whose receive ready field is set to indicate a second specific round robin order in a group formed by P number out of N 20 transmission links, P being a positive integer smaller than N and finally sending each ATM data cell in said series of ATM data cells in said second specific round robin order.

Brief Description of the Drawings

25 For a more complete understanding of the invention and for further objects and advantages thereof, reference may now be made to the following description, taken in conjunction with the accompanying drawings, in which:

Figure 1 shows a sample configuration involving AIM UNIs;

30 Figure 2 shows a sample configuration involving AIM BICIs or PNNIs;

Figure 3 depicts the ATM cell structure defined in the ITU Recommendation;

35 Figure 4 is a schematic illustration of multiplexing and de-multiplexing of ATM cells over AIMs and links;

Figure 5 shows preassigned cell header values for use by the physical layer;

Figure 6 shows a header pattern of an AIM OAM cell, according to one embodiment of the invention;

Figure 7 indicates allocation of OAM functions in the information field according to one embodiment of the invention;

5 Figure 8 is a schematic illustration of multiplexing and demultiplexing of AIM OAM cells according to one embodiment of the invention; and

10 Figure 9 is a schematic illustration of periodic multiplexing and demultiplexing of AIM OAM cells during transmission of ATM data cells according to one embodiment of the invention.

Figure 10 is an ICP cell format used for implementing embodiments of the invention.

15 Figure 11 shows a circumstance where node A may be using a group of three links to send data to node B and another group of two links to send data to node C.

Figure 12 shows a typical sequence of cells on a three link group.

Figures 13a and 13b show an algorithmic flow chart for the start-up procedure according to one embodiment.

Figure 14 indicates how Figures 13a and 13b should be connected.

20 Figure 15 shows a configuration when links with Link ID=1 are added to groups with active Link ID=0 links.

Figure 16 shows the timing diagram of the test procedure of the invention.

Figure 17 illustrates looping of a test pattern at node B.

25 Figure 18 depicts diagrammatically a scenario where the test pattern procedure can be used to detect a link that is not connected to the expected IMA.

Figure 19 shows cases (1,...,7) when the ICP cell preceding the stuff event and the stuff ICP cells are corrupted.

30 Detailed Description of Preferred Embodiments

There are requirements that have to be considered when defining a new ATM inverse multiplexing protocol. These requirements are:

35 • It must multiplex and demultiplex an ATM cell stream distributed in a round robin manner over multiple links such as T1/E1 links.

- It must adjust up to 32 milliseconds of differential link delays between individual links in a case where T1/E1 links are used.
- It must reconfigure multiple links in the event that a link has to be added, deleted or is considered inadequate to provide service.
- 5 • It must be defined for not only UNIs but also for PNNIs and BICIs.
- It must be transparent to the devices handling the convergence of ATM cells into the PDH signal.
- It must be transparent to the devices dealing with ATM layer cells.

10 The present invention achieves all of the above requirements and solves the problems discussed earlier. The invention relates to a new ATM inverse multiplexing scheme that makes use of a physical layer operation administration and maintenance (OAM) cell which has been properly defined. This OAM cell is defined to contain valuable information to allow proper operation of the ATM inverse
15 multiplexing mechanism and also to provide opportunity for handling a link failure situation.

The newly defined OAM cell is called an AIM OAM cell or AIM Sequence Number (SN) cell and is mainly designed to carry a cell sequence number and a feedback link status field. The sequence number in the SN cell is made available for the receiver end for recovering the initial cell stream from the incoming links. The feedback link status is made available for the receiver to inform the transmitter, by sending its own SN cells, that it is receiving cells and it is also an integral part of the same round robin mechanism, that is to say, the feedback link status value SN cells in either direction must agree with each other, although the sequence numbers at both ends may be different. When the receiver sends its own SN cells in response to the transmitter that it is receiving cells, it is in fact acknowledging that the receiver is ready to receive subsequent ATM data cells.

30 Figure 4 shows how the ATM cells are multiplexed and then demultiplexed over AIMs in one direction. At the transmitting node, an AIM 10 takes a series of ATM cells from an ATM layer device. It spreads ATM cells and transmits each cell over each of N transmission links, N being a positive integer. The order of transmission is in round robin fashion. This process is called an inverse multiplexing. At the receiving node, cells from N links are inverse demultiplexed (assembled) and sent to an ATM layer device by an AIM 14. The same

order must be employed at this node to recover a proper sequence of cells. Both nodes must be aware of the round robin order which is to be employed. Upon initialization, therefore, both AIMs send a series of SN cells in round robin fashion over the links, e.g. T1/E1 links. This 5 allows the receiver AIM at both nodes to establish the sequence in which to read cells from the incoming links as well as to adjust relative link delay.

As mentioned above, the SN cell also carries an extra field used by each T1/E1 link to indicate that both AIMs belong to the same link 10 round robin. This information can then be used locally to determine if at the corresponding local node a link should be added, removed or maintained in the round robin.

When a change of link configuration occurs by a link being added, removed or declared as being down, each node sends a series of 15 SN cells to allow the far-end node to reestablish the sequence of cells to read from the incoming links.

The protocol of the invention calls for a physical layer OAM cell that is defined to be exclusively processed by the ATM inverse multiplexers. The new cell structure has to be consistent with the cell 20 structure defined in ITU Recommendation I.361. Figure 3 shows that structure of UNI/NNI ATM cells that is defined in I.361. ITU Recommendations I.361 and I.432 state that ATM cells consisting of a header whose VPI and VCI fields are set to zero are reserved for use by the physical layer. So far three preassigned values of the cell header are 25 reserved for use by the physical layer. They are shown in Figure 5.

The physical layer OAM cell according to one embodiment of the invention is then defined by using a non-assigned value by setting the PT field to "111" or some such code. This is shown in Figure 6.

The AIM OAM cell payload is then available for exchanging 30 information between AIMs. The cell payload consists of a series of fields whose locations are shown in Figure 7.

The following fields are identified for the AIM protocol according to one embodiment of the invention:

- ATM inverse multiplexing remote defect indicator (AIMRDI) – one 35 octet is allocated and the proposed coding is all "1".
- ATM inverse multiplexing far-end receiver ready (AIMFERR) – one octet is allocated and the proposed coding is all "1".

- ATM inverse multiplexing cell sequence number (AIMCSN) – it is defined to contain the sequence number of the cells sent over the multiple links handled by the ATM inverse multiplexers. It is designed so as to have a sufficiently large cycle to allow the ATM inverse multiplexer to absorb link delays of up to 32 milliseconds.
- 5 There are 16 bits allocated to the AIMCSN field. The counting is then done modulo 65536.
- Cell error control (CEC) – is used to detect errors in the cell payload. It is proposed to CRC-10 as proposed in ITU Recommendation I.432.
- 10 • Reserve field (R) – contains the octet pattern of "01101010", which is the same as that of the idle cell as proposed in ITU Recommendation I.361.

Referring to Figure 8, the ATM inverse multiplexing protocol according to one embodiment is described in detail below. The figure 15 only shows one direction for clarity. It is shown that transmission links are all T1/E1 links but, of course, links can have a different speed from T1/E1, as long as all the links have the same speed. A transmitting node 20 collects digital data consisting of a series of ATM data cells from the ATM layer. An ATM inverse multiplexer (AIM) at 20 the transmitting node spreads over multiple links 24 for transmission to a receiving node 26. An AIM at the receiving node reassembles ATM data cells received from the links in proper order and sends them to the ATM layer.

25 Link Start-up

Upon connection start-up, AIMs at both nodes start inserting AIM OAM cells (AIM SN cells) carrying cell sequence number over the available links in round robin fashion. The sequencing of cells is based on the order in which cells have to be transmitted on the virtual link 30 (composed of N physical links). However, the sequence number is only carried over the SN cells. This sequence number assignment allows the receiving AIM to retrieve the original cell sequence. The receiving node queues the received AIM SN cells until it determines the sequence in which to read the ATM data cells from the incoming links 35 as well as the differential delay among individual links. Then, it starts sending AIM SN cells of its own, with the AIMFERR field set to "one" for each link which is now considered "ready" to receive ATM traffic.

From that moment, the receiving node knows the sequence of cells coming from the links. A link is being considered available if cells are currently delineated and AIM-RDI is not received on the incoming link.

5 When the transmitting node starts sending the AIM SN cells, it starts a time-out of 100 milliseconds within which the receiving node determines the sequence of AIM SN cells which it is receiving from the links. The time-out expires unless the transmitting node receives the "ready" signal from all the links which were considered available at the
10 initialization.

15 If there is no "ready" link when the time-out expires, the local node reevaluates the availability of the links (using cell delineation), starts sending AIM SN cells over the available links and re-starts the time-out. This procedure is repeated until at least one available link is declared ready.

20 When the time-out has been cancelled due to the reception of the "ready" signal from all the available links, or when there is at least one link ready when the time-out expires, the local end starts sending ATM layer cells over the "ready" link or links using the same round robin order used at the time of initialization.

25 Each ATM layer cell gets a sequence number assigned to it, but only the AIM SN cells carry that number across the links. After both nodes have started sending ATM layer cells, they periodically send a series of "n" AIM SN cells over links to allow the receiving node to readjust the differential delays among the links. The value of "n" is equal to the number of "ready" links used to carry ATM layer cells. This means that the receiving node always scans for AIM SN cells in order to determine if it is always reading cells from the multiple links in the correct order. How often a burst of "n" AIM SN cells is sent
30 depends upon the link utilization but the maximum period has been set to 50 milliseconds in one embodiment. An example of cells sent over a virtual link composed of a plurality of transmission links is shown in Figure 9.

35 Link Reconfiguration

As mentioned above, the protocol according to the invention

also deals with link reconfiguration. There are three possible cases in which reconfiguration can occur:

- a new link has to be added to the round robin;
- a link has to be intentionally removed from the round robin; and
- 5 • a link has been declared inadequate to provide service (e.g. link down by failure etc.).

In the first two cases, a node starts link reconfiguration by performing the same initialization process for a connection start-up, that is to say, it chooses a round robin order among the links and starts 10 sending AIM SN cells using the chosen round robin order. The receiving node, seeing the occurrence of AIM SN cells, will stop transmitting traffic and starts sending AIM SN cells of its own while adjusting itself to receive traffic from the incoming links. The rest of the protocol is as described earlier for start-up.

15 Two possible cases for reporting failure conditions are:

1) A medium problem reported through Loss of Delineation (LCD) failure condition. In this case, the corresponding link should not be used for service. When detecting LCD, SN cells with AIMRDI set to "1" should be sent over the corresponding outgoing link. SN cells 20 containing AIM-RDI do not carry a valid sequence number. At the far-end, the detection of AIM-RDI signals will indicate that the link is not to be used. Therefore, if LCD or AIM-RDI is detected, the link is no longer considered available and won't be allowed to be part of the 25 round robin on both sides. When this happens, SN cells are re-sent over the remaining links to re-establish the connection between the two ends.

2) Cells are lost without an LCD or AIM-RDI being reported. For instance, this would occur when a few cells are discarded by the physical layer device due to bit errors in cell headers. This would cause 30 the cell sequencing to be affected. One symptom would be the detection of SN cells whose number is no longer the same as that expected (since one or more of the previous cells are missing). In this case, the problem might be partially corrected by re-adjusting the receiver buffering system (assuming the difference between the SN cell number 35 and the expected number is small). The other symptom would be the absence of a SN cell on one link when getting a burst of SN cells on all the other links. In that case, the local end would have to force the links

to be re-configured. If the symptoms described above reoccur over a given period, the bad link may have to be removed from the round robin.

5 A link reconfiguration also occurs when it takes too much time for one link to receive cells from the far-end (receiving) node, that is to say, no cells within, e.g., 32 milliseconds.

Cell Sequence Number Range

10 Because an ATM inverse multiplexer must absorb a differential delay between individual links of up to 32 milliseconds in one embodiment, and because the system needs to deal with a maximum of 8 T1/E1 links, it is necessary to have a sequence number whose modulo is large enough to accommodate such delay.

15 As a practical example, the following parameters for E1 are considered:

- full rates: 2.048 Mbit/s
- payload rates: $30/32 \times 2.048 \text{ Mbit/s} = 1.92 \text{ Mbps}$
- ATM cell: 53 bytes
- ATM cell period time: 221 microseconds
- 20 • ATM cells/32 milliseconds = 144 cells.

For a delay of up to 32 milliseconds on each link, there is a need to queue cells for at least the same length of time on each link. Therefore, in this example, this means that a delay of up to a period of 144 cells between two links is possible. An ATM inverse multiplexer 25 can handle a maximum of 8 T1/E1 links. Therefore up to 1008 cells (144*7 cells) must be queued at one time by one ATM inverse multiplexer. This requires a sequence number modulo large enough to cover this scenario. A simple case in one embodiment is to use a 16-bit count that has modulo 65536.

30 The ATM inverse multiplexing protocol according to the invention realizes the following characteristics:

- it is applicable to UNIs, BICIs and PNNIs (any applicable points in an ATM network);
- it does not affect the ATM cell header of currently defined cells;
- 35 • it does not require a change to current ATM physical layer devices dealing with the convergence of ATM cells into T1/E1;
- it operates transparently to the ATM layer;

- it does not require a change to current devices dealing with the processing of the ATM layer cell;
- it is self-configuring among available links upon start-up and self-reconfiguring among the multiple links in the case when a new link has to be added, deleted or considered inadequate to provide service; and
- it includes a sequence number whose modulo is large enough to meet the requirement of a large differential delay among the links.

10 According to an embodiment of the present invention, special Inverse mux Controller Processor (ICP) cells are created. The ICP cell format is shown in Figure 10. In this embodiment, ATM cells are transmitted over an N number of links, N being a positive integer. Cell ID is set for ICP cell and Link ID identifies links being used. As 15 seen in Figure 11, a node A may be using a group of three links to send data to node B and another group of two links to send data to node C. Nodes A and B form one IMA group and nodes A and C form another. IMA groups are identified by IMA ID (Tx and Rx IMA ID). ATM cells on each link are grouped in a certain number (e.g., M) of cells to form 20 IMA frames which belong to substantially same time scale. An IMA frame sequence number field counts cells in each group. Figure 12 shows a typical sequence of cells on a three link group. One ICP cell is sent at a set location in each group for each link. M can be any number, for example 32, 64, 128 and 256. In the figure, on link 0, the ICP cells 25 have their cell offset set to zero (i.e., they are the first cell in the IMA frame). On link 1, the ICP cells have the ICP cell offset set to 3 and on link 2, the ICP cells have their ICP cell offset set to 1. In practice, these ICP cells should be distributed even more over the IMA frame but are shown closer for ease of illustration. The ICP cell also includes fields 30 concerning with stuffing action and test control action, both of which will be described in detail below.

35 A status and control change indication field is used to indicate an update of the link status field value. The value is incremented to indicate a change on at least one of the link status fields. The field is also used as a tag to differentiate link status changes over time. The field will always remain set to the same value as long as there will be no change on any link status fields. If one or more link status fields

need to be modified, ICP cells with new link status values incremented will be sent over all the links.

A group status and control field is used to indicate status of the group of links at a connection start-up, link addition and abort procedures. In particular, the start-up procedure would become complicated if one of the links does not meet all the criteria to be part of the group (e.g. link cleared of defects, corresponding incoming/outgoing links in loss of delay synchronization). This requires intervention of the operator to remove the bad link(s) from the link group. Figures 13a and 13b as combined in the way shown in Figure 14 show an algorithmic flow chart for the start-up procedure according to one embodiment. Instead of requiring all N number of links in the group being good before starting up, the procedure of the invention proceeds as long as there is at least P out of N "good" link available for service. A good link is defined as a link that is not exhibiting link and loss-of-frame defects, recognized as being part of the group, and having an acceptable link differential delay. The value of P could be any number less than N . The figure also shows cases where the group start-up is aborted and later resumed.

The ICP format also includes a field for link's connectivity testing at the time of start-up, link addition or link re-activation. In particular, a protocol must always ensure proper connectivity of the links which belong to a group. For example, Figure 15 shows when links with Link ID=1 are added to groups with active Link ID=0 links. Two new links (Link ID=1 links) are inversely connected with respect to the expected configuration. The difficulty is then that both Link ID and IMA ID are the same for two independent IMAs. This can cause an invalid configuration which will not be detected. This kind of problem can occur at a start-up, link addition or even at a link re-activation. For example, it is possible to reactivate a link which has exhibited a link failure since it has been accidentally inverted with another link with or without the same Link ID and IMA ID. IMA ID is also used to detect loopback situations. This problem occurs when both end nodes of the IMA virtual link wants to use the same IMA ID. The invention addresses this problem also.

According to a further embodiment, the above problems are solved by sending a test pattern contained in ICP cells over a link to be

validated. The test pattern will be looped back over all the other links in the group at the far end node. This ensures that the tested link is connected to the same end node as the other links.

When one IMA node wants to determine if one or more links are connected to the same far end IMA node, it selects one link for testing. Figure 16 shows the timing diagram of the test procedure of the invention. The transmit node sets the Tx Test control field for link test in an ICP cell. It identifies Link ID, Tx IMA ID, Rx IMA ID and inserts Test Pattern in the cell. It continues to send same ICPs. The receive end node receives the ICPs and copies Tx Test Pattern onto Rx Test Pattern. It keeps sending same test pattern copied on Rx Test Pattern field over its all outgoing links. Figure 17 illustrates looping of a test pattern at node B. The transmit node receives and verifies Rx Test Pattern returned over the other links. It now knows that all the links which belong to the same IMA group.

Figure 18 depicts diagrammatically a scenario where the test pattern procedure can be used to detect a link that is not connected to the expected IMA. Links identified as Link ID=1 are being added to existing groups using links with Link ID=0. Link 30 carries the IMA B Test Pattern to be looped back by IMA A and links 32 and 34 carry the IMA D Test Pattern to be looped back to IMA B. In this case, the wrong test pattern will be received by IMA B. If IMA B was not commanding a Test Pattern loopback, IMA B would simply not receive the right test pattern.

The Test Pattern procedure also allows to deal with some pathological cases. One of them is when two IMA connected to the same far end IMA node are trying to start-up at the same time. The far end (receive end) node will have to determine which set of links (corresponding to one end) it wants to be connected to. This will require the received end to select the link(s) to be part of the group. as mentioned above, the receive end shall only respond to one Test Pattern command at a time that has been validated over one or more links that are or likely to be recognized as part of the group.

In accordance with a further embodiment, a stuff cell is inserted to control cell rate decoupling between the links, in order to accommodate the use of links non-synchronized to each other within the link groups. The transmitting node may be locked to one clock

source or may be plesiochronous. When plesiochronous, one of the buffers at the transmitting node may be depleted. To prevent underrun, the cell stuffing procedure is invoked. When there is one clock source, the buffer will never deplete. The transmitting node send 5 ICP cells which indicate a cell is stuffed at a certain location within the IMA frame. Any cell can be used for stuffing as long as the location is indicated so that the receiving node can remove it. In actual embodiment, the transmitting node repeats the ICP cell containing the stuff code indicating that "this cell is 1 out of 2 stuff cells". The 10 receiving node uses the stuff indications over the ICP cells to determine when to remove stuff cells from the incoming cell stream. The receiving node relies on at least one ICP cell with a correct CRC-10. A more robust approach is to look for a majority of valid codes.

Figure 19 shows cases (1,...,7) when the ICP cell preceding the 15 stuff event and the stuff ICP cells are corrupted. Corrupted cells are indicated by crosses in the figure. SICP indicate stuffing control cells. The receiving node maintains synchronization for cases 1, 2, and 3. The receiving node optionally maintain synchronization for cases 4, 5, and 6. the receiving node optionally maintain synchronization for case 20 7 if $b > 2$ and when passing stuff indication over more than one of the previous ICP cells. "b" is the number of invalid/corrupted ICP cells before declaring the link OIF (out-of-IMA frame).

WHAT IS CLAIMED IS:

1. In an inverse multiplexing digital data over a connection consisting of a plurality of transmission links, said data containing a series of ATM data cells, a method of testing connectivity of the transmission links comprising steps of:
 - 1 sending a series of inverse multiplexing control cells containing a test message over one of the transmission links;
 - 2 receiving a series of inverse multiplexing control cells containing received test messages over all the remaining transmission links; and
 - 3 verifying the received test messages to determine the connectivity of the transmission links.
- 15 2. The method of testing connectivity of the transmission links according to claim 1 wherein the step of sending comprising further steps of:
 - 1 setting a field within each of the series of inverse multiplexing control cells to identify a transmission link to be tested, and;
- 20 3 sending the series of inverse multiplexing control cells containing a preset test pattern over the transmission links.
3. The method of testing connectivity of the transmission links according to claim 2, comprising further steps of:
 - 1 receiving a series of inverse multiplexing control cells to see if the test pattern are copied thereon, and
 - 2 sending the series of inverse multiplexing control cells over all the remaining transmission links.
- 30 4. The method of testing connectivity of the transmission links according to claim 1, comprising further steps of:
 - 1 setting a timeout period during which a series of inverse multiplexing control cells are continuously sent over the transmission links.



5. The method of testing connectivity of the transmission links according to claim 1, wherein the step of verifying comprising further steps of:


monitoring the test message on the received inverse
5 multiplexing control cells to determine which transmission links
belong to one group.

6. A method of inverse multiplexing digital data over a connection consisting of a plurality of transmission links, said data 10 containing a series of ATM data cells, comprising steps of:
sending a series of inverse multiplexing control cells indicating a specific round robin order in which the series of ATM data cells are to be transmitted over the connection;
receiving from the plurality of transmission links a series of 15 inverse multiplexing control cells whose receive ready field is set;
sending each ATM data cell in said series of ATM data cells in said specific round robin order; and
further sending two consecutive inverse multiplexing control cells in a frame, indicating cell stuffing.

20

7. In an inverse multiplexing digital data over a connection consisting of a plurality of transmission links, said data containing a series of ATM data cells, an apparatus for testing connectivity of the transmission links comprising:
25 a transmitter for sending a series of inverse multiplexing control cells containing a test message over one of the transmission links;
a receiver for receiving a series of inverse multiplexing control cells containing received test message over all the remaining transmission links; and
30 a message control device for verifying the received test message to determine the connectivity of the transmission links.





GFC: General Flow Control
VPI: Virtual Path Identifier
VCI: Virtual Channel Identifier
PT: Payload Type
CLP: Cell Loss Priority
HEC: Header Error Check

Fig 3




Fig 4




Fig 8

Fig 5

Cell Purpose	Octet 1	Octet 2	Octet 3	Octet 4	Octet 5
Idle cell identification	00000000	00000000	00000000	00000001	01010010
Cell-based interface F1 OAM cell	00000000	00000000	00000000	00000011	01011100
Cell-based interface F3 OAM cell	00000000	00000000	00000000	00001001	01101000

Fig 6

Octet 1	Octet 2	Octet 3	Octet 4	Octet 5 (HEC)
00000000	00000000	00000000	00001111	11001011

OCTET	ALLOCATION
1	AIM-RDI
2	AIMFERR
3	AIMCSN(MSB)
4	AIMCSN(LSB)
5	R
6	R
7	R
8	R
9	R
10	R
11	R
12	R
13	R
14	R
15	R
16	R
17	R
18	R
19	R
20	R
21	R
22	R
23	R
45	R
46	R
47	R
48	CFC (10)

FIG 7

MSB: Most Significant Byte

LSB: Least Significant Byte




Fig 9

OCTET	FIELD
1-5	ATM Cell Header
6	IMA Label
7	Cell ID, Link ID
8	IMA Frame Sequence Number
9	ICP Cell Offset
10	Link Stuff Indication
11	Status & Control Change Indication
12	Tx IMA ID
13	Rx IMA ID
14	Group Status & Control
15	Tx Test Control
16	Tx Test Pattern
17	Rx Test Pattern
18	Link 0 Information
19	Link 1 Information
20-49	Link 2-31 Information
50-51	Unused
52-53	CRC-10

FIG 10

Fig 11





Fig 12






Fig 13 a

*Fig 13 b**Fig 14*




Fig 15




Fig 17




Fig 18




Fig 16




Fig 19

INTERNATIONAL SEARCH REPORT

International Application No

PCT/CA 97/00610

A. CLASSIFICATION OF SUBJECT MATTER
IPC 6 H04Q11/04

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
IPC 6 H04Q

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	WO 96 08120 A (STRATACOM INC) 14 March 1996	1,7
X	see page 5	6
A	WO 96 17489 A (NORTHERN TELECOM LTD) 6 June 1996	1,7
X	see page 4, line 22 - line 31	6
	see page 7, line 1 - line 6	
A	EP 0 584 398 A (SIEMENS AG) 2 March 1994 see column 1, line 46 - column 2, line 6	1,6,7

Further documents are listed in the continuation of box C.

Patent family members are listed in annex.

* Special categories of cited documents :

- "A" document defining the general state of the art which is not considered to be of particular relevance
- "E" earlier document but published on or after the international filing date
- "L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
- "O" document referring to an oral disclosure, use, exhibition or other means
- "P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.

"8" document member of the same patent family

1

Date of the actual completion of the international search

13 November 1997

Date of mailing of the international search report

08/12/1997

Name and mailing address of the ISA

European Patent Office, P.O. 5818 Patentlaan 2
NL - 2280 MV Rijswijk
Tel. (+31-70) 340-2040, Tx. 31 651 epo nl,
Fax: (+31-70) 340-3016

Authorized officer

Veen, G

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
WO 9608120 A	14-03-96	US 5617417 A AU 3321495 A CA 2199383 A EP 0780046 A	01-04-97 27-03-96 14-03-96 25-06-97
WO 9617489 A	06-06-96	US 5608733 A CA 2204171 A EP 0795259 A	04-03-97 06-06-96 17-09-97
EP 0584398 A	02-03-94	JP 6177907 A US 5394398 A	24-06-94 28-02-95

